arXiv:1607.05527v1 [cs.CG] 19 Jul 2016

An Approximation Algorithm for the Art Gallery Problem*

Édouard Bonnet[†1] and Tillmann Miltzow[‡1]

[1]Institute for Computer Science and Control Hungarian Academy of Sciences (MTA SZTAKI)

Abstract

Given a simple polygon $\mathcal{P}$ on n vertices, two points x, y in $\mathcal{P}$ are said to be visible to each other if the line segment between x and y is contained in $\mathcal{P}$. The Point Guard Art Gallery problem asks for a minimum set S such that every point in $\mathcal{P}$ is visible from a point in S. The set S is referred to as guards. Assuming integer coordinates and a specific general position assumption, we present the first $O(\log \mathrm{OPT})$-approximation algorithm for the point guard problem.[1] This algorithm combines ideas of a paper of Efrat and Har-Peled [16] and Deshpande et al. [13, 14]. We also point out a mistake in the latter.

1 Introduction

Given a simple polygon $\mathcal{P}$ on n vertices, two points x, y in $\mathcal{P}$ are said to be visible to each other if the line segment between x and y is contained in $\mathcal{P}$. The point-guard art gallery problem asks for a minimum set S such that every point in $\mathcal{P}$ is visible from a point in S. The set S is referred to as guards.

A huge amount of research is committed to the studies of combinatorial and algorithmic aspects of the art gallery problem, such as reflected by the following surveys [19, 30, 32, 33]. Most of this research, however, is not focused directly on the art gallery problem but on variants, based on different definitions of visibility, restricted classes of polygons, different shapes and positions of guards, and so on. The most natural definition of visibility is arguably the one we gave above. Other possible definitions are: x sees y if the axis-parallel rectangle spanned by x and y is contained in $\mathcal{P}$; x sees y if the line segment $\mathrm{seg}(x, y)$ intersects $\mathcal{P}$ at most c times, for some value of c; x sees y if there exists a straight-line path from x to y within $\mathcal{P}$ with at most c bends. Common shapes of polygons include: simple polygons, polygons with holes, simple orthogonal polygons, x-monotone polygons and star-shaped polygons. Common placements of guards include: vertex guards and point guards as defined above, but also edge-guard (guards are edges of the polygon), segment guards (guards are interior segments of the polygon) and perimeter guards (guards must be placed on the boundary of $\mathcal{P}$). The art gallery problem variants, also distinguish the way that the polygon is covered. For example, it might be required that every point is seen by two different guards; sometimes it might be required that every point is covered by one guard of each color; and recently the community gets also interested in conflict-free guard colorings. That is, every point has a color such that it sees

*supported by the ERC grant PARAMTIGHT: "Parameterized complexity and the search for tight complexity results", no. 280152.

[†]edouard.bonnet@lamsade.dauphine.fr

[‡]t.miltzow@gmail.com

[1]For the benefit of the reviewers, we uploaded on youtube a video that describes informally the main ideas of the paper: `https://youtu.be/k5CDeimSuBM`. The video is meant as supplementary material and is not required for anything that is coming.

exactly one guard of that color. In 1978, Steve Fisk proved elegantly that $\lfloor n/3 \rfloor$ guards are always sufficient and sometimes necessary for a polygon with n vertices [18]. Five years earlier, Victor Klee has posed this question to Václav Chvátal, who soon gave a more complicated solution. This constitutes the first combinatorial result related to the art gallery problem.

On the algorithmic side, very few variants are known to be solvable in polynomial time [15, 29] and most results are on approximating the minimum number of guards [13, 14, 16, 20, 25, 26]. Many of the approximation algorithms are based on the fact that the range space defined by the visibility regions has bounded VC-dimension for simple polygons [21, 22, 34]. This makes it easy to use the algorithmic ideas of Clarkson [8, 12].

On the lower bound side, Eidenbenz et al. [17] showed NP-hardness and inapproximability for most relevant variants. In particular, they show for the main variants that there is no c-approximation algorithm for simple polygons, for some constant c. For polygons with holes, they can even show that there is no $o(\log n)$-approximation algorithm. Also, their reduction from Set-Cover implies that the art gallery problem is W[2]-hard on polygons with holes and that there is no $n^{o(k)}$ algorithm, to determine if k guards are sufficient, under the Exponential Time Hypothesis [17, Sec.4]. Recently, a similar result was shown for simple polygons (without holes) [6, 7].

Despite, the large amount of research on the art gallery problem, there is only one exact algorithmic result on the point guard variant. The result is not so well-known and attributed to Micha Sharir [16]: One can find in $n^{O(k)}$ time a set of k guards for the point guard variant, if it exists. This result is quite easy to achieve with standard tools from real algebraic geometry [3] and apparently quite hopeless to prove without this powerful machinery (see [4] for the very restricted case $k = 2$). Despite the fact that the algorithm uses remarkably sophisticated tools, it uses almost no problem-specific insights and no better exact algorithms are known. Some recent ETH-based lower bounds [6, 7] suggest that there might be no better exact algorithm even for simple polygons.

Regarding *approximation* algorithms for the point guard variant, the results are similarly sparse. For general polygons, Deshpande et al. gave a randomized pseudo-polynomial time $O(\log n)$-approximation algorithm [13, 14]. However, we show that their algorithm is not correct. Efrat and Har-Peled gave a randomized polynomial time $O(\log |OPT_{\text{grid}}|)$-approximation algorithm by restricting guards to a very fine grid [16]. However, they could not prove that their S_{grid} grid solution is indeed an approximation of an optimal guard placement. Developing the ideas of Deshpande et al. in combination of the algorithm of Efrat and Har-Peled, we attain the first randomized polynomial-time approximation algorithm for simple polygons. Here, OPT denotes an optimal set of guards and OPT_{grid} an optimal set of guards that is restricted to some grid. At last, we want to mention that there exist approximation algorithms for monotone and rectilinear polygons [27], when the very restrictive structure of the polygon is exploited.

To understand the lack of progress, note that the art gallery problem can be seen as a geometric hitting set problem. In a hitting set problem, we are given a universe U and a set of subsets $S \subseteq 2^U$ and we are asked to find a smallest set $X \subseteq U$ such that $\forall r \in S\ \exists x \in X : x \in r$. Usually the set system is given explicitly or can be at least easily restricted to a set of polynomial size. In our case, the universe is the entire polygon (not just the boundary) and the set system is the set of visibility regions (given a point $x \in \mathcal{P}$, the visibility region $\text{Vis}(x)$ is defined as the set of points visible from x). The crucial point is that the set system is *infinite* and no one has found a way to restrict the universe to a finite set (see [1, 11] for some attempts). We also wish to quote a recent remark by Bhattiprolu and Har-Peled [5]:

"One of the more interesting versions of the geometric hitting set problem, is the art gallery problem, where one is given a simple polygon in the plane, and one has to select a set of points (inside or on the boundary of the polygon) that "see" the whole polygon. While much research has gone into variants of this problem [30], nothing is known as far as an approximation algorithm (for the general problem). The difficulty arises from the underlying set system being

infinite, see [16] for some efforts in better understanding this problem."

Here, we present the first approximation algorithm for simple polygons under some mild assumptions.

Assumption 1 (integer vertex representation)**.** Vertices are given by integers, represented in binary.

An *extension* of a polygon $\mathcal{P}$ is a line that goes through two vertices of $\mathcal{P}$.

Assumption 2 (general position assumption)**.** No three extensions meet in a point of $\mathcal{P}$ which is *not* a vertex and no three vertices are collinear.

Note that we allow that three (or more) extensions meet in a vertex or outside the polygon.

Theorem 1. *Under Assumptions 1 and 2, there is a randomized $O(\log |OPT|)$-approximation algorithm for* Point Guard Art Gallery *for simple polygons that runs in polynomial time in the size of the input.*

The main technical idea is to show the following lemma:

Lemma 2 (Global Visibility Containment)**.** *Let $\mathcal{P}$ be some (not necessarily simple) polygon. Under Assumptions 1 and 2, it holds that there exists a grid Γ and a guard set $S_{grid} \subseteq \Gamma$, which sees the entire polygon and $|S_{grid}| = O(|S|)$, where S is an optimal guard set.*

To be a bit more precise, let M be the largest appearing integer. Then the number of points in Γ is polynomial in M. This is potentially exponential in the size of the input. Thus algorithms that rely on storing all points of Γ explicitly do not have polynomial worst case running time. The algorithm of Efrat and Har-Peled [16] does *not* store every point of Γ explicitly and, with the lemma above, the algorithm gives an $O(\log |OPT|)$-approximation on the grid Γ.

While Lemma 2 tells us that we can restrict our attention to a finite grid, when considering constant factor approximation, the same is not known for exact computation. In particular, it is not known whether the Point Guard problem lies in NP. Recently, some researchers popularized an interesting class, called $\exists\mathbb{R}$, being somewhere between NP and PSPACE [9, 10, 28, 31]. Many geometric problems, for which membership in NP is uncertain, have been shown to be complete for this class. This suggests that there might be indeed no polynomial sized witness for these problems as this would imply $NP = \exists\mathbb{R}$. The history of the art gallery problem suggests the possibility that the Point Guard problem is $\exists\mathbb{R}$-complete. If $NP \neq \exists\mathbb{R}$, then this would imply that there is indeed no hope to find a witness of polynomial size for the Point Guard problem.

In computational geometry and discrete geometry many papers assume that *no three points lie on a line*. Often this assumption is a pure technicality, in other cases the result might indeed be wrong without this assumption. In our case, we do believe that Lemma 2 could be proven without Assumption 2, but it seems that some new ideas would be needed for a rigorous proof. See [2] for an example where the main result is that some general position assumption can be weakened. The idea of general position assumptions is that a small random perturbation of the point set yields the assumption with probability almost 1. In case that the points are given by integers small random perturbations, destroy the integer property. But random perturbations could be performed in a different way, by first multiplying all coordinates by some large constant $2^C \in \mathbb{N}$ and then add a random integer x with $-C \leqslant x \leqslant C$.

The integer representation assumption (Assumption 1) seems to be very strong as it gives us useful distance bounds not just between any two different vertices of the polygon, but also between any two objects that do not share a point (see Lemma 5). On the other hand, real computers work with binary numbers and cannot compute real numbers with arbitrary precision.

The real-RAM model was introduced as a convenient theoretical framework to simplify the analysis of algorithms with numerical and/or geometrical flavors, see for instance [24, page 1] and [23, Remark 23.1]. Also note that Assumption 1 can be replaced by assuming that all coordinates are represented by rational numbers with specified nominator and denominator. (There could be other potentially more compact ways to specify rational numbers.) Multiplying all numbers with the smallest common multiple of the denominators takes polynomial time, makes all numbers integers and does not change the geometry of the problem.

Given a polygon $\mathcal{P}$, we will always assume that all its vertices are given by *positive* integers in binary. (This can be achieved in polynomial time.) We denote by M the largest appearing integer and we denote by $\operatorname{diam}(\mathcal{P})$ the largest distance between any two points in $\mathcal{P}$. Note that $\operatorname{diam}(\mathcal{P}) < 2M$. We denote $L = 20M > 10$. Note that $\log L$ is linear in the input size. We define the grid

$$\Gamma = (L^{-11} \cdot \mathbb{Z}^2) \cap \mathcal{P}.$$

Note that all vertices of $\mathcal{P}$ have integer coordinates and thus are included in Γ.

Theorem 3 (Efrat, Har-Peled [16])**.** *Given a simple polygon $\mathcal{P}$ with n vertices, one can spread a grid Γ inside $\mathcal{P}$, and compute an $O(\log OPT_{\text{grid}})$-approximation for the smallest subset of Γ that sees $\mathcal{P}$. The expected running time of the algorithm is*

$$O(n\, OPT_{\text{grid}}^2 \log OPT_{grid}\, \log(n\, OPT_{\text{grid}})\, \log^2 \Delta),$$

where Δ is the ratio between the diameter of the polygon and the grid size.

The term OPT_{grid} refers to the optimum, when restricted to the grid Γ. For the solution S that is output by the algorithm of Efrat and Har-Peled holds $|S| = O(|OPT_{\text{grid}}| \log |OPT_{\text{grid}}|)$. However, Efrat and Har-Peled make no claim on the relation between $|S|$ and the actual optimum $|OPT|$. Note that the grid size equals $w = L^{-11}$, thus $\Delta \leqslant L^{11+1} = L^{12}$ and consequently $\log \Delta \leqslant 12 \log L$, which is polynomial in the size of the input.

Efrat and Har-Peled implicitly use the real-RAM as model of computation: elementary computations are expected to take $O(1)$ time and coordinates of points are given by real numbers. As we assume that coordinates are given by integers, the word-RAM or integer-RAM is a more appropriate model of computation. All we need to know about this model is that we can upper bound the time for elementary computations by a polynomial in the bit length of the involved numbers. Thus, going from the real-RAM to the word-RAM only adds a polynomial factor in the running time of the algorithm of Efrat and Har-Peled. Therefore, from the discussion above we see that it is sufficient to prove Lemma 2.

Organization. In Section 2, we will describe the counterexample to the algorithm of Deshpande *et. al.* This will be also very useful as a starting point of Section 3, where we will give an elaborate overview of the forthcoming proofs. Detailed and formal proofs are presented at Section 4, where they can be read in logical order without references to Section 3. Finally in Section 5, we briefly indicate the remaining open questions.

2 Counterexample

In this section, we will point out a mistake in the algorithm of Deshpande et al. [13, 14]. This mistake though constitutes an interesting starting point for our purpose.

The algorithm by Deshpande et al. can be described from a high level perspective as follows: maintain and refine a triangulation T of the polygon until for every triangle $\Delta \in T$ holds the so called *local visibility containment* property. The local visibility containment property of Δ certifies that every point $p \in \Delta$ can only see points that are also seen by the vertices of Δ.

However, we will argue that it is impossible to attain the local visibility containment property with any finite triangulation; hence, the algorithm never stops.

Again, we want to mention that the paper of Deshpande et al. has ideas that helped to achieve the result of the present paper. In particular, we will show that the local visibility containment property does indeed hold most of the time.

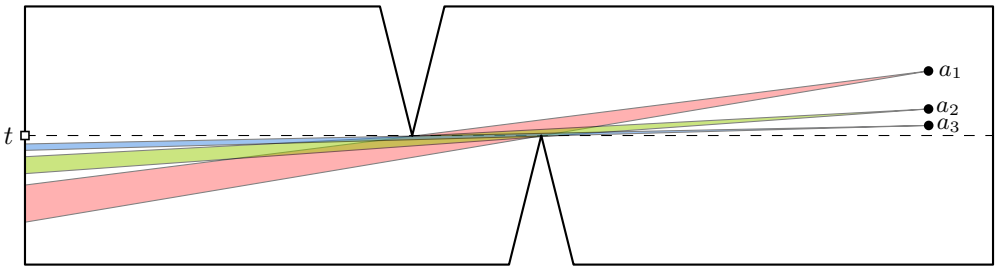

Figure 1: Illustration of the counterexample to the algorithm of Deshpande et al. [13, 14]

Example 4 (Refutation of Deshpande et al. [13, 14])**.** See Figure 1, for the following description. We have two opposite reflex vertices with supporting line ℓ. The points $(a_i)_{i\in\mathbb{N}}$ are chosen closer and closer to ℓ on the right side of the polygon. None of the a_i's can see t, as this would require to be actually on ℓ. Further, it is easy to see that the *visible interval* next to t gets smaller and smaller as well. In fact, we can choose the points $(a_i)_{i\in\mathbb{N}}$ in a way that their intervals will be all disjoint.

Consider now any finite collection of points C in the vicinity of the $(a_i)_{i\in\mathbb{N}}$. We will show that there is some a_i, which sees some interval close to t, that is not seen by any point in C. Recall that no point sees the entire interval around t, but the visibility of the a_i's come arbitrarily close to t. Thus, there is some a_i that sees something that is not visible by any of the points of C. □

3 Proof Overview

In this section, we will describe all the proof ideas without going into too much technical details. Definitions are given by figures and some technical conditions are not stated at all. The reader who feels uncomfortable with this is deferred to Section 4, where all definitions are rigorously made and all lemmas are presented with formal, detailed proofs.

Some readers might also be concerned with the wrong proportions of the figures. Some distances which are supposedly very small are displayed fairly large and vice versa. The reader has to keep in mind that all figures only indicate principal behavior.

Our high level proof idea is that the local visibility property holds for every point x that is far enough away from all extension lines. (Recall that the extension of two vertices is the line that contains these vertices.) In a second step, we will show that it is impossible to be close to more than 2 extensions at the same time. We will add one vertex for each extension that x is close two. Recall that vertices are also in Γ.

It turns out that the first step is considerably more tedious than the second one. The reason for that is that many of the elementary steps are not true in the naive way, one might think of at first. Thus we have to carefully define the sense in which they are true and handle the other case in a different manner. The tricky bit is usually to identify these cases and to carefully define them. All proofs are elementary otherwise.

3.1 Benefit of integer coordinates

The integer coordinate assumption implies not just that the distance between any two vertices is at least 1 but it also gives useful lower bounds on distances between any two objects of interest that do not share a point. For example the distance between an extension ℓ and a vertex v not on ℓ is at least L^{-1}. Also the angle between any two non-parallel extensions is at least L^{-2}. (Recall that L is an upper bound on the diameter and the largest appearing integer.) As these bounds are important for the intuition of the forthcoming ideas, we will proof one of them.

Claim. *Let $\ell = \ell(w_1, w_2)$ be the extension of the two vertices w_1, w_2 of $\mathcal{P}$ and $v \notin \ell$ some other vertex of $\mathcal{P}$. Then, it holds $d = dist(v, \ell) \geqslant L^{-1}$.*

Here $\text{dist}(v, \ell)$ denotes the euclidean distance between v and ℓ.

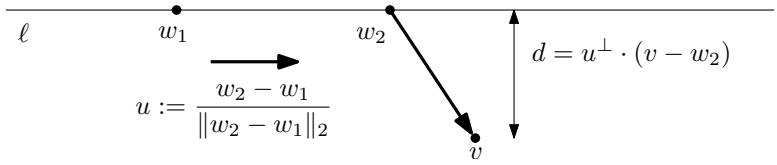

Figure 2: Computing the distance between a line and a vertex.

Proof. The distance d can be computed as

$$d = \frac{|(v - w_1) \cdot (w_2 - w_1)^\perp|}{\|w_2 - w_1\|} \geqslant \frac{1}{\text{diam}(\mathcal{P})} \geqslant \frac{1}{L}.$$

See Figure 2 for a way to derive this elementary formula. Here $\cdot$ denotes the scalar product and $x^\perp$ is the vector x rotated by 90° and $\|x\|_2$ is the euclidean norm of x.

The key inside is that the nominator of this formula is at least 1 as it is a non-zero integer by assumption. The denominator is upper bounded by the diameter of $\mathcal{P}$, which is in turn upper bounded by L. □

All other lower bounds are derived in the same spirit, however with worth bounds. As we choose our grid width smaller than any of these bounds, we will be in the very fortunate situation that everything looks very simple from a local perspective.

3.2 Surrounding Grid Points

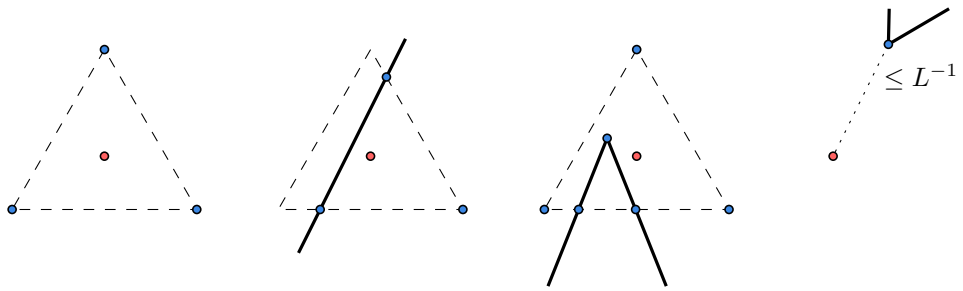

Figure 3: The red point indicates a point of the original optimal solution. The blue points indicate the surrounding grid points that we choose. The polygon is indicated by bold lines. From left to right, we have the three cases: interior case, boundary case and corner case. To the very right, we indicate that in every case vertices of $\mathcal{P}$ with distance less than L^{-1} are also included in α-grid$^*(x)$.

Given a point $x \in OPT$ we define α-grid(x) as some grid points around x, as displayed in Figure 3. The parameter $\alpha \ll 1$ determines how far away these points are. As the grid-width w is much smaller than α, we have a choice on how exactly to place these points. See Section 4.3 for a precise definition. If there is any vertex v of $\mathcal{P}$ with $\text{dist}(x, v) \leqslant L^{-1}$, then this point will be included into α-grid$^*(x)$.

3.3 Local Visibility Containment

For any extension ℓ we define an s-bad region, see the gray area in Figure 4a for an illustration. Note that the bad region consists of two connected components, each being a triangle. The parameter $s = \tan(\beta)$ as indicated in the figure. Further, for each point x, the visibility region can be decomposed into triangles as indicated in Figure 4b.

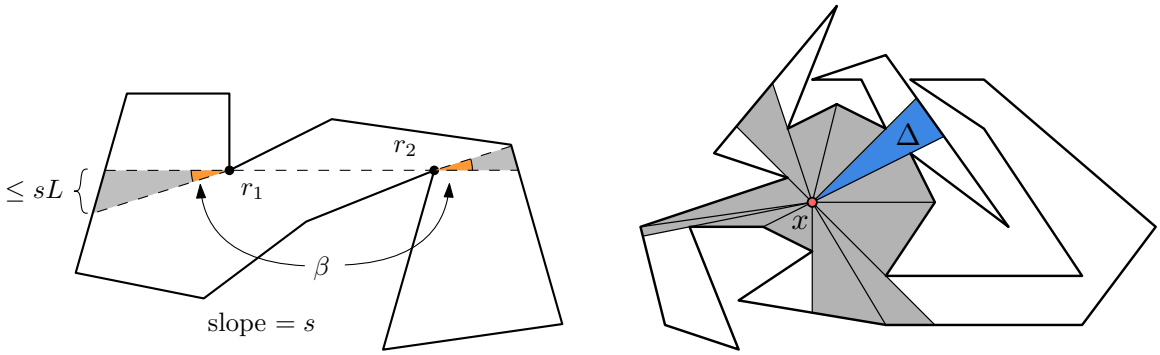

(a) A polygon with two opposite reflex vertices and their s-bad region.

(b) The star triangle decomposition of the visibility region of x.

Figure 4

Let Δ be some triangle of the visibility region of x (in blue in Figure 4b). Then the main lemma asserts that α-grid$^*(x)$ sees Δ except, if x is in an s-bad region of the vertices defining Δ, see Lemma 11. (With the right choice of α and s.) Important is the one-to-one correspondence between the triangles that cannot be seen and the extension line that we can make responsible for it.

The first technical lemma concerns the possibility that there is a reflex vertex q that blocks the visibility of the points of α-grid(x) onto Δ, see Figure 5. We can show that this can happen only in a negligible amount. For this recall that α-grid(x) is very close to x, but q must have distance at least L^{-1}, as otherwise it would be included into α-grid$^*(x)$. We call this phenomenon *limited blocking*.

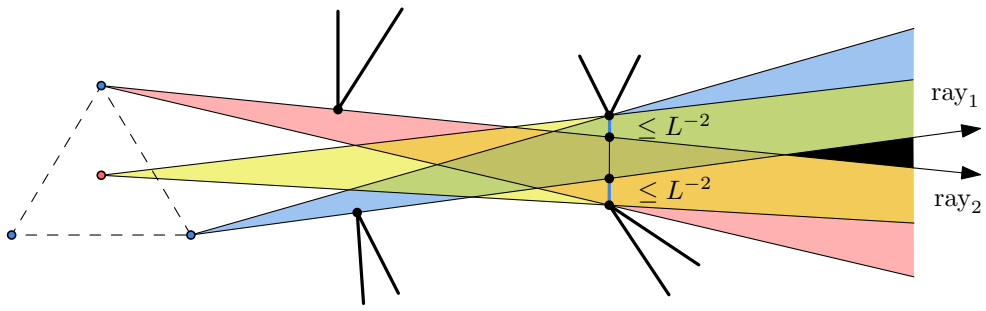

Figure 5: The visibility of the grid points $g \in \alpha$-grid(x) can be blocked, but we can bound the amount by which it is blocked. The key idea to show that R_2 can be seen by α-grid(x) is to show that the region indicated in solid black is empty.

Further the triangle Δ is split into a *small triangle* (R_1) and a trapezoid(R_2), as indicated

in Figure 6a. We show separately, for R_1 and R_2 that α-grid$^*(x)$ sees these two regions.

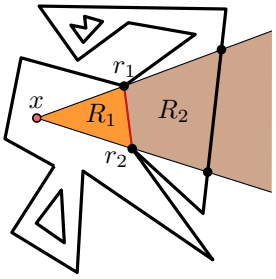

(a) To show that each triangle of the visibility region is visible by α-grid$^*(x)$, we treat the small triangle R_1 and the trapezoid R_2 individually. In particular, as we do not make use of the finiteness of R_2, we just assume it is an infinite cone.

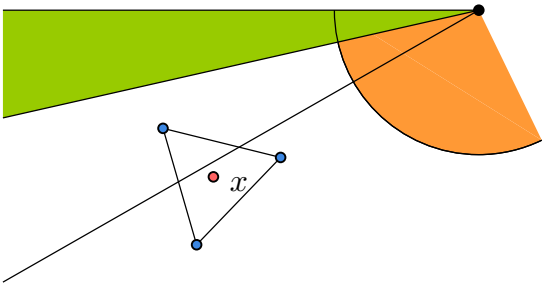

(b) The point x is not in an s-bad region implies that α-grid(x) is not in an $s/2$ bad region.

Figure 6

To prove that α-grid$^*(x)$ sees R_1 is not so difficult. We have already argued that reflex vertices v (with $\text{dist}(x, v) \geqslant L^{-1}$) can only block a small part of the visibility of α-grid(x). (Recall that any vertex v with $\text{dist}(x, v) < L^{-1}$ is included in α-grid$^*(x)$.) With some simple case distinctions, we can argue that this visibility is sufficient to see R_1. In particular, the argument does not rely on x being outside a bad region.

To prove that R_2 can be seen by α-grid$^*(x)$ is more demanding. As it seems not useful to use the boundedness of R_2, we just assume it to be an infinite cone and we try to show that α-grid$^*(x)$ sees this cone. Obviously, the part of $\partial\mathcal{P}$ "behind" $\text{seg}(r_1, r_2)$ is not considered blocking. The crucial step to show that R_2 can be seen by α-grid$^*(x)$ is to show that the black region as indicated in Figure 5 does not exist. The idea is that this is implied if ray_1 and ray_2 diverge. In other words if ray_1 and ray_2 never meet and the black region is empty. For this purpose, we make use of the fact that $\text{dist}(g_1, g_2) \approx \alpha$, for any $g_1, g_2 \in \alpha$-grid(x) by definition, while $\text{dist}(r_1, r_2) \geqslant 1$, because of integer coordinates. Thus intuitively, the distance of ray_1 and ray_2 is closer at its apex than at the segment $\text{seg}(r_1, r_2)$. For the proof to gp through we further need α-grid$^*(x)$ not to be in an $s/2$-bad region, which follows from x not being in an s-bad region, see Figure 6b. The proof becomes technically even more tedious as we have to take into consideration that the visibility of α-grid$^*(x)$ might be partially blocked. This also forces us to introduce embiggened bad regions, which differ only marginally from bad regions as defined above.

3.4 Global Visibility Containment

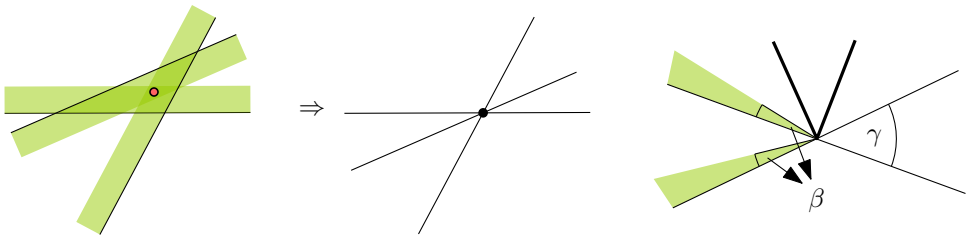

Figure 7: Three bad regions meeting in an interior point implies that the extensions must meet in a single point. No two bad regions intersect in the vicinity of a vertex, as they are defined by some angle $\beta \ll L^{-2}$. But the angle γ between any two extensions is at least L^{-2}.

Given a minimum solution OPT, we describe a set $G \subseteq \Gamma$ of size $O(|OPT|)$. We will also

show that G sees the entire polygon. For each $x \in OPT$, G contains α-grid$^*(x)$. Further if x is contained in an s-bad region, G contains at least one of the vertices defining it. It is clear by the previous discussion that G sees the entire polygon, as the only part that is not seen by α-grid$^*(x)$ are some small regions, which are entirely seen by the vertices bounding it.

It remains to show that there is no point in three bad regions. For this, we heavily rely on the integer coordinates and the general position assumption. Note that the integer coordinate assumption implies not just that the distance between any two vertices is at least 1 but also that the distance between any extension ℓ and a vertex v not on ℓ is at least L^{-1}. Also the angle between any two extensions is at least L^{-2}. (Recall that L is an upper bound on the diameter and the largest appearing integer.) These bounds and other bounds of this kind imply that if any three bad regions meet in the interior, then their extension lines must meet in a single point, see Figure 7. We exclude this by our general position assumption. Close to a vertex, we use a different argument: No two bad regions intersect in the vicinity of a vertex, as bad regions are defined by some angle β with $\tan(\beta) \ll L^{-2}$. But the angle γ between any two extensions is at least L^{-2}.

This implies that each x is in at most 2 bad regions and $|G| = (7+2)|OPT| = O(|OPT|)$.

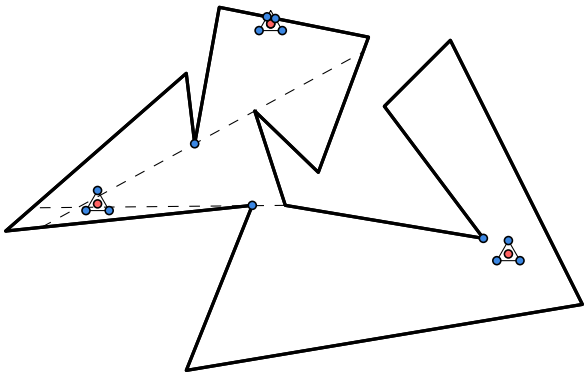

Figure 8: The red dots indicate the optimal solution. The blue dots indicate the The red dot on the top is in the interior case and four grid points are added around it. The red dot on the left is too close to two supporting lines and we add one of the reflex vertices of each of the supporting lines. The red dot to the right has distance less than L^{-1} to a reflex vertex, so we add that vertex as well to G.

4 Proofs

4.1 Preliminaries

Polygons and visibility. For any two distinct points v and w in the plane, we denote by $\mathrm{seg}(v,w)$ the segment whose two endpoints are v and w, by $\mathrm{ray}(v,w)$ the ray starting at v and passing through w, by $\ell(v,w)$ the supporting line passing through v and w. We think of $\ell(v,w)$ as directed from v to w. Given a *directed* line ℓ, we denote by ℓ^+ the half-plane to the right of ℓ bounded by ℓ. We denote by ℓ^- the half-plane to the left of ℓ bounded by ℓ. We also denote by $\mathrm{disk}(v,r)$ the disk centered in point v and whose radius is r, and by $\mathrm{dist}(a,b)$ the distance between point a and point b. We denote by $\mathrm{diam}(\mathcal{P})$ the diameter of $\mathcal{P}$ and by M the largest appearing integer of any vertex of $\mathcal{P}$. We further assume that all coordinates are represented by positive integers. (As already stated earlier, this can be achieved in polynomial time.) We define $L = 20M$. This implies $\mathrm{diam}(\mathcal{P}) \leqslant 2M$ and $10\,\mathrm{diam}(\mathcal{P}) \leqslant L$. A polygon is *simple* if it is not self-crossing and has no holes. For any point x in a polygon $\mathcal{P}$, $\mathrm{Vis}(x)$, denotes the *visibility region* of x within $\mathcal{P}$, that is the set of all the points $y \in \mathcal{P}$ such that segment $\mathrm{seg}(x,y)$ is entirely contained in $\mathcal{P}$.

symbol	definition
$\mathcal{P}$	denotes underlying polygon (we assume all coordinates to be positive)
$\mathrm{Vis}(x)$	visibility region of x
$\mathrm{diam}(\mathcal{P})$	the largest distance between any two points in $\mathcal{P}$
M	largest appearing integer of $\mathcal{P}$
L	$= 20M$ (it holds $\mathrm{diam}(\mathcal{P}) \leqslant 10L$.)
Γ	the grid $(L^{-11} \cdot \mathbb{Z}^2) \cap \mathcal{P}$.
$\mathrm{seg}(v, w)$	segment with endpoints v and w
$v(p, q) \in S^1$	direction from p to q
$\mathrm{ray}(p, q)$	p, q points in the plane: ray with apex p in direction $v(p, q)$
$\mathrm{ray}(p, v)$	p points in the plane, $v \in S^1$ direction: ray with apex p in direction v
$\ell(v, w)$	line through v and w directed from v to w
ℓ^+, ℓ^-	the half-plane to the right respectively to the left of line ℓ
$\mathrm{disk}(v, r)$	disk centered in point v with radius r
$\mathrm{dist}(a, b)$	euclidean distance between a and b.
$\mathrm{dist}_{\mathrm{horizontal}}(a, b)$	$\lvert a_x - b_x \rvert$
$\mathrm{dist}_{\mathrm{vertical}}(a, b)$	$\lvert a_y - b_y \rvert$
α-grid (x)	some grid points around x, see Figure 11
α-grid *(x)	α-grid(x) and including a possible vertex v with $\mathrm{dist}(r, x) \leqslant L^{-1}$
$\mathrm{cone}(x, r_1, r_2)$	a cone with apex x bounded by $\mathrm{ray}(x, r_1)$ and $\mathrm{ray}(x, r_2)$
$\mathrm{cone}(x)$	we use this notation, when r_1 and r_2 are clear from context.
$\mathrm{apex}(x, y)$	see Figure 18b and Definition 15

Figure 9: This is a glossary of all used notation. Not all of which is introduced in the Preliminaries.

4.2 Benefit of Integer Coordinates

The way we use the integer coordinate assumption is to infer distance lower bounds between various objects of interest. The proofs are technical and not very enlightening. Fortunately they are short.

Lemma 5 (Distances)**.** *Let $\mathcal{P}$ be a polygon with integer coordinates and L as defined above. Let v and w be vertices of $\mathcal{P}$, ℓ and ℓ' supporting lines of two vertices, and p and q intersections of supporting lines.*

1. *dist*$(v, w) > 0 \Rightarrow$ *dist*$(v, w) \geqslant 1$.
2. *dist*$(v, \ell) > 0 \Rightarrow$ *dist*$(v, \ell) \geqslant L-1$.
3. *dist*$(p, \ell) > 0 \Rightarrow$ *dist*$(p, \ell) \geqslant L^{-5}$.
4. *dist*$(p, q) > 0 \Rightarrow$ *dist*$(p, q) \geqslant L^{-4}$.
5. *Let $\ell \neq \ell'$ be parallel. Then dist*$(\ell, \ell') \geqslant L^{-1}$.
6. *Let $\ell \neq \ell'$ be any two non-parallel supporting lines and α the smaller angle between them. Then holds* $\tan(\alpha) \geqslant 8L^{-2}$.
7. *Let $a \in \mathcal{P}$ be a point and ℓ_1 and ℓ_2 be some non-parallel lines with dist*$(\ell_i, a) < d$*, for $i = 1, 2$. Then ℓ_1 and ℓ_2 intersect in a point p with dist*$(a, p) \leqslant dL^2$.

Proof. Case 1 seems trivial, but follows the same general principal as the other bounds. All these distances, are realized by geometric objects, and these geometric objects are represented with the help of the input integers. In order to compute theses distances some elementary calculations

are performed and solutions can be expressed as fractions of the input integers. Using the lower bound one and the upper bound L or $L/2$ on these integers or derived expressions give the desired results.

Consider Case 2. Let ℓ be the supporting line of w_1 and w_2. The distance d between v and ℓ can be expressed as $\frac{|(v-w_1)\cdot(w_2-w_1)^\perp|}{\|w_2-w_1\|}$, where $\binom{x_1}{x_2}^\perp = \binom{x_2}{-x_1}$ represents the orthogonal vector to x and $\cdot$ indicates the dot product. As v is not on ℓ the nominator is lower bounded by one. The denominator is upper bounded by the diameter $\operatorname{diam}(\mathcal{P})$.

Consider Case 3. Let p be the intersection of $\ell_1 = \ell(v_1, v_2)$ and $\ell_2 = \ell(u_1, u_2)$. Then p is the unique solution to the linear equations $(u_2 - u_1)\cdot(p - u_1) = 0$ and $(v_2 - v_1)\cdot(p - v_1) = 0$. So let us write this as an abstract linear equation $A \cdot x = b$. By Cramer's rule holds $x_i = \frac{\det(A_i)}{\det(A)}$, for $i = 1, 2$. Here $\det(.)$ is the determinant and A_i is the matrix with the i-th column replaced by b. It is easy to see that $det(A)$ is bounded by L^2. Let $\ell = \ell(w_1, w_2)$ be a different line. Observe that the points w_1, w_2 and p lie on a grid Γ^* of width $1/\det(A)$. After scaling everything with $det(A)$ the diameter $\operatorname{diam}(\mathcal{P})$ of this grid becomes $det(A)L \leqslant L^3$. Thus by Case 2 the distance between the line ℓ and the grid point p is lower bounded by $(1/L^*) \geqslant 1/L^3$. Scaling everything back by $1/det(A) \geqslant L^{-2}$ yields the claimed bound.

Consider case 4. Let p be the intersection of $\ell_1 = \ell(v_1, v_2)$ and $\ell_2 = \ell(u_1, u_2)$. And likewise let q be the intersection of $\ell_3 = \ell(w_1, w_2)$ and $\ell_2 = \ell(t_1, t_2)$. Then p is the unique solution to the linear equations $(u_2 - u_1)\cdot(p - u_1) = 0$ and $(v_2 - v_1)\cdot(p - v_1) = 0$. So let us write this as an abstract linear equation $A \cdot x = b$. By Cramer's rule holds $x_i = \frac{det(A_i)}{det(A)}$, for $i = 1, 2$. Thus the coefficients of $p - q$ can be represented by $\frac{a}{det(A)\cdot det(A')}$ for some value $a \neq 0$ and some matrices A and A'. Note that $det(A)$ and $det(A')$ are bounded by L^2 from above.

Case 5 follows from Case 2. Let $\ell' = \ell(w_1, w_2)$. Then $\operatorname{dist}(\ell, \ell') = \operatorname{dist}(\ell, w_1) \geqslant 1/L$.

Consider case 6. We assume $\ell = \ell(v_1, v_2)$ and $\ell' = \ell(w_1, w_2)$. The expressions $\tan(\alpha)$ can be computed by

$$\frac{\|(v_2 - v_1) \times (w_2 - w_1)\|}{\|(v_2 - v_1)\| \cdot \|(w_2 - w_1)\|} \geqslant \frac{1}{\operatorname{diam}(\mathcal{P})^2} \geqslant \frac{8}{L^2}.$$

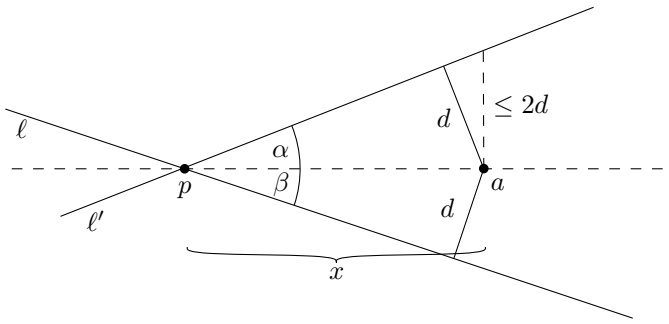

Figure 10

For case 7, see Figure 10 and the notation therein. Without loss of generality we assume $\alpha \geqslant \beta$. It follows from elementary calculus that $z \leqslant \tan(z) \leqslant 2z$, for any sufficiently small z. This implies $8L^{-2} \leqslant \tan(\alpha + \beta) \leqslant 2(\alpha + \beta) \leqslant 4\alpha \leqslant 4\tan\alpha$. We follow from $x \tan\alpha \leqslant 2d$, it holds $x \leqslant 2d/\tan\alpha \leqslant dL^2$. □

4.3 Defining surrounding grid points

The point of this section is to define for a point x some grid-points α-grid(x) surrounding x. Morally, α-grid(x) should see whatever x sees. However, this cannot be achieved in general.

Definition 6 (Rounding)**.** Given a point $x \in \mathcal{P}$, we define the point $\operatorname{round}(x) = g \in \Gamma \subset \mathcal{P}$ as the closest grid-point to x. In case that there are several grid-points with the same minimum

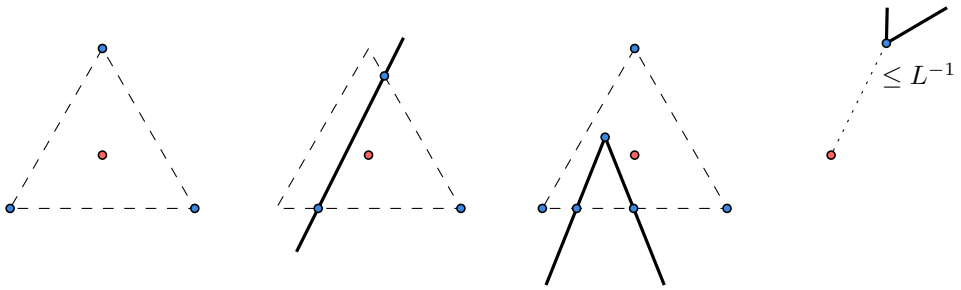

Figure 11: The red point indicates a point of the original optimal solution. The blue points indicate the surrounding grid points that we choose. The polygon is indicated by bold lines. From left to right, we have the three cases: interior case, boundary case and corner case. To the very right, we indicate that in every case vertices of $\mathcal{P}$ with distance less than L^{-1} are also included in α-grid$^*(x)$.

distance to x, we choose the one with lexicographic smallest coordinates. The important point here is that round(x) is uniquely defined.

Definition 7 (Surrounding Grid points)**.** Given a point $x \in \mathcal{P}$ and a number $\alpha \in [0, L^{-2}]$, we define α-grid(x) as a set of grid points around x. The definition depends on the position of x and the value α. Let c be a circle with radius alpha and center x. Then there exists a unique triangle $\Delta(x)$ inscribed c such that the lower side of $\Delta(x)$ is horizontal. We distinguish three cases. In the *interior* case, $\Delta(x)$ and $\partial\mathcal{P}$ are disjoint. In the *boundary* case, $\Delta(x)$ and $\partial\mathcal{P}$ have a non-empty intersection, but no vertex of $\mathcal{P}$ is contained in $\Delta(x)$. In the *corner* case, on vertex of $\mathcal{P}$ is contained in Δ. It is easy to see that this covers all the cases. We also say a point x in the interior case, and so on. In the *interior case* α-grid(x) is defined as follows. Let v_1, v_2, v_3 be the vertices of Δ. Then the grid points round(v_i), for all $i = 1, 2, 3$ form the surrounding grid points. In the *boundary case* α-grid(x) is defined as follows. Let S be the set of vertices of Δ and all intersection points of $\partial\mathcal{P}$ with $\partial\Delta(x)$. Then the grid points round(v), for all $v \in S \cap \mathcal{P}$ form the surrounding grid points. In the *corner case* α-grid(x) is defined as follows. Let S be the set of vertices of Δ. all intersection points of $\partial\mathcal{P}$ with $\partial\Delta$ and the vertex of $\mathcal{P}$ contained in Δ. Then the grid points round(v), for all $v \in S \cap \mathcal{P}$ form the surrounding grid points. In any case, if there is a reflex vertex r with $\text{dist}(x, r) \leqslant L^{-1}$ then we include r in the set α-grid$^*(x) = r \cup \alpha$-grid(x) as well. We will usually denote the points in α-grid(x) with g_1, g_2 or just g.

As we will choose $\alpha \gg w$, the difference between x and round(x) is negligible and thus we will assume that $x = \text{round}(x)$.

4.4 Local Visibility Containment

Definition 8 (Local Visibility Containment property)**.** We say a point x has the *α-local visibility containment property* if

$$\text{Vis}(x) \subseteq \bigcup_{g \in \alpha\text{-grid}^*(x)} \text{Vis}(g).$$

Definition 9 (Opposite reflex vertices and bad regions)**.** Given a polygon $\mathcal{P}$ and two reflex vertices r_1 and r_2, consider the supporting line $\ell = \ell(r_1, r_2)$ restricted to $\mathcal{P}$. We say r_1 is *opposite* to r_2 if both incident edges of r_1 lie on the opposite side of ℓ as the edges of r_2, see Figure 12a, for an illustration. Given two opposite reflex vertices r_1 and r_2, we define their *s-bad region* as the union of the two triangles as in Figure 12a, where we set the slope to s. Alternatively, let β be the angle at r_i, then we define $\tan(\beta) = s$.

We denote by p_i the point on $\text{seg}(r_1, r_2)$ with $\text{dist}(p_i, r_i) = L^{-2}$, for $i = 1, 2$. We define the *embiggened s-bad region* of r_1 and r_2 as the union of the two triangles Δ_1 and Δ_2, one of which is displayed in Figure 12b. Formally triangle Δ_i has one corner p_i. The two edges incident to p_i form a slope of s. One of these edges is part of $\ell(r_1, r_2)$ the other is in the interior of $\mathcal{P}$. The last remaining edge is part of $\partial\mathcal{P}$.

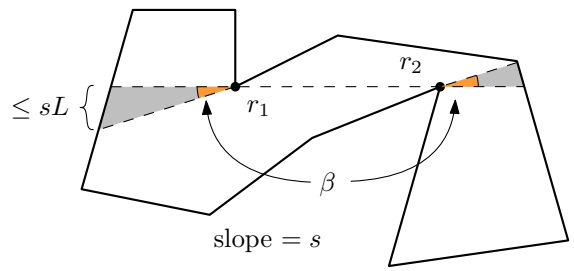

(a) A polygon with two opposite reflex vertices and their s-bad regions.

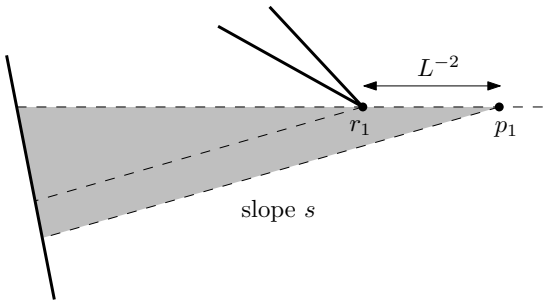

(b) The embiggened s-bad region.

Figure 12: Illustration of bad regions and embiggened bad regions.

Definition 10 (Triangle decomposition of visibility regions and cones)**.** Given a polygon $\mathcal{P}$ and some point x, we define the *star triangle decomposition* of $\text{Vis}(x)$ as follows. Let $v_1, \ldots, v_t$ be the vertices visible from x in clockwise order. Then any two consecutive vertices u, v together with x define a triangle Δ of $\text{Vis}(x)$. Note that u, v are not necessarily the vertices of Δ, see Figure 13a. Also note that the definition does not require the polygon to be simple. (The polygon could have holes.)

We denote by $\text{cone}(x, u, v) = \text{cone}(x)$ the cone with apex x that is bounded by $\text{ray}(x, u)$ and $\text{ray}(x, v)$, see Figure 13b. We will assume that u and v are implicitly known and omit to mention them if there is no ambiguity. Note that $\text{cone}(x)$ is unbounded and not contained in $\mathcal{P}$.

Let u, v be vertices, not necessarily visible from some other point g. Let $s \subseteq \text{seg}(u, v)$, the part of the segment that is visible from g. Then we define

$$\text{cone}(g) = \bigcup_{t \in s} \text{ray}(g, t).$$

See the blue cone in Figure 13b. We say $\alpha\text{-grid}(x)$ *sees* $\text{cone}(x)$ if

$$\text{cone}(x) \subseteq \bigcup_{g \in \alpha\text{-grid}(x)} \text{Vis}(g) \quad \cup \quad \bigcup_{g \in \alpha\text{-grid}(x)} \text{cone}(g).$$

Note that $\text{cone}(g)$ is not contained into $\text{Vis}(g)$ as the cone is unbounded, whereas the visibility region is contained in $\mathcal{P}$. We define $\alpha\text{-grid}^*(x)$ *sees* $\text{cone}(x)$ in the same fashion.

The purpose of this Section is to prove Lemma 11.

Lemma 11 (Special Local Visibility Containment Property)**.** *Let r_1 and r_2 be two consecutive vertices in the clockwise order of the vertices visible from $x \in \mathcal{P}$ and let x be outside the s-bad region of the vertices r_1 and r_2. We make the following assumptions:*

1. $s \leqslant L^{-3}$
2. $\alpha \leqslant L^{-7}$.
3. $16L\alpha \leqslant s$

Then α-grid$^*(x)$ sees cone(x).

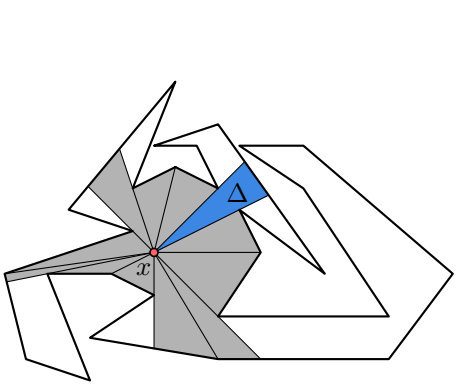

(a) The star triangle decomposition of the visibility region of x.

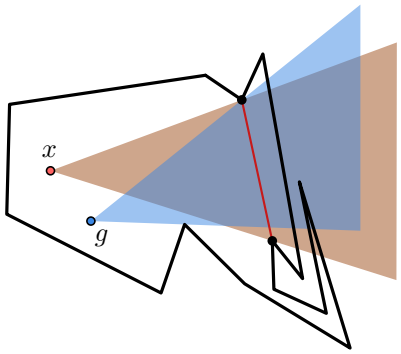

(b) The cone of x and g are displayed. Both cones are unbounded. The cone of g is not bounded by the vertices that define it.

Figure 13: Illustration of the star triangle decomposition and cones.

Note that, we do not exclude that x is outside *any* s-bad region but only outside the bad region of r_1 and r_2. This subtlety makes the proof more complicated, as it might be that x is in an s-bad region with respect to some other pair of reflex vertices. For instance q, r_1, could be such a pair. It is not difficult to construct an example such that q blocks part of the visibility of α-grid(x). The proof idea is to show that this blocking is very limited (see Lemma 13) and then use only this partially obstructed visibility.

Seeing cone(x) implies obviously also that the triangle $\Delta \subseteq$ cone(x) is seen by α-grid(x). An immediate consequence is the following nice lemma.

Lemma 12 (Local Visibility Containment Property)**.** *Every point $x \in \mathcal{P}$ outside any s-bad region has the α-local visibility containment property. We make the following assumptions:*

1. $s \leqslant L^{-3}$
2. $\alpha \leqslant L^{-7}$.
3. $16\, L\, \alpha \leqslant s$

Proof. By Lemma 11, every triangle Δ defined by the star triangle decomposition of Vis(x) is seen by α-grid$^*(x)$. □

For Lemma 13, see Figure 14a for an illustration of the setup. The next Lemma deals with the very special situation that there exists a reflex vertex q blocking the visibility of some $g \in \alpha$-grid(x), although $q \notin$ cone(x). The critical assumption is dist$(q, x) \geqslant L^{-1}$.

Lemma 13 (Limited Blocking)**.** *Let $x \in \mathcal{P}$, let r_1, r_2 be two vertices and $g \in \alpha$-grid(x) such that $g \notin$ cone(x) and $\alpha \leqslant L^{-7}$. Further, let $q \in$ cone(g) be a reflex vertex with dist$(x, q) > L^{-1}$. We denote by p the intersection point between $\ell(g, q)$ and seg(r_1, r_2). Let g and r_2 be on the same side of cone(x). Then holds:*

$$d = \mathit{dist}(p, r_2) \leqslant L^{-2}.$$

Proof. We show first dist$(g, \ell(r_1, r_2)) \geqslant L^{-1}$. Let ℓ be the line parallel to $\ell(r_1, r_2)$ containing the point q. Clearly, these two lines have distance at least L^{-1}, by Lemma 5 Item 2. And thus also g has at least this distance to $\ell(r_1, r_2)$, see Figure 14a.

For the other part of the proof let us assume that $\ell(r_1, r_2)$ is vertical. We denote by c the *vertical* distance between g and $\ell(q, r_2)$ and by s the point realizing this distance, see Figure 14a. Note that dist$(g, q) \geqslant$ dist$(x, q) -$ dist$(x, g) \geqslant L^{-1} - \alpha \geqslant 1/(2L)$. By Thales Theorem we know:

$$\frac{d}{\mathrm{dist}(p, q)} = \frac{c}{\mathrm{dist}(g, q)} \leqslant \frac{c}{1/(2L)} = 2cL.$$

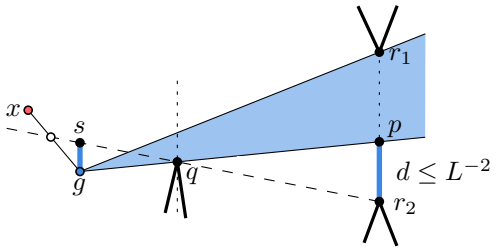

(a) The reflex vertex q blocks the visibility of g. Fortunately, we can show that only a very small part of the visibility is actually blocked.

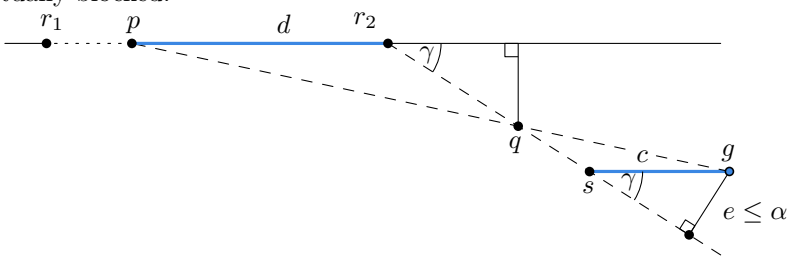

(b) This Figure illustrates the proof of Lemma 13.

Figure 14

This implies

$$d \leqslant 2cL \operatorname{dist}(p,q) \leqslant 2cL^2$$

It is sufficient to show $c \leqslant \alpha L^2$, as $\alpha \leqslant L^{-7}$. For this purpose, we rotate $\ell(r_1, r_2)$ again so that it becomes horizontal, see Figure 14b. Now c is the vertical distance between g and $\ell(q, r_2)$. We define $e = \operatorname{dist}(g, \ell(r_2, q))$ and we denote by γ the angle at r_2 formed by $\ell(r_2, q)$ and $\ell(r_1, r_2)$. We consider here only the case that $\gamma < 90°$, the other case is symmetric. We note that γ appears again at s. This implies

$$\frac{e}{c} = \sin(\gamma).$$

It is also easy to see that

$$\sin(\gamma) = \frac{\operatorname{dist}(\ell(r_1, r_2), q)}{\operatorname{dist}(r_2, q)} \geqslant \frac{L^{-1}}{L} = L^{-2}.$$

Further note that $\ell(r_2, q)$ must intersect $\operatorname{seg}(g, x)$ as q blocks the visibility of g but not of x. Thus $e = \operatorname{dist}(g, \ell(r_2, q)) \leqslant \alpha$. This implies

$$c = \frac{e}{\sin(\gamma)} \leqslant \alpha L^2.$$

And thus it holds

$$d \leqslant 2cL^2 \leqslant 2(\alpha L^2)L^2 = \alpha L^5 \leqslant L^{-7}L^5 = L^{-2}.$$ □

Lemma 14 (Small Triangle)**.** *Let $x \in \mathcal{P}$ and let r_1 and r_2 be two vertices such that $\Delta = \Delta(x, r_1, r_2) \subseteq \mathcal{P}$. Then it holds that $\Delta(x, r_1, r_2)$ is seen by α-grid$^*(x)$, with $\alpha \leqslant L^{-7}$.*

Proof. We consider first the case that $\operatorname{dist}(x, \ell(r_1, r_2)) \leqslant L^{-1}/2$. Note that there is at least one $g \in \alpha\text{-grid}(x)$ on the same side of $\ell(r_1, r_2)$ as x. It follows by Lemma 5 that there can be no vertex of $\mathcal{P}$ blocking the visibility of g to $\Delta(x, r_1, r_2)$. Thus from now on, we can assume that $\operatorname{dist}(x, \ell(r_1, r_2)) \leqslant L^{-1}/2$ and this implies in particular that x and $\alpha\text{-grid}(x)$ are on the same side of $\ell(r_1, r_2)$.

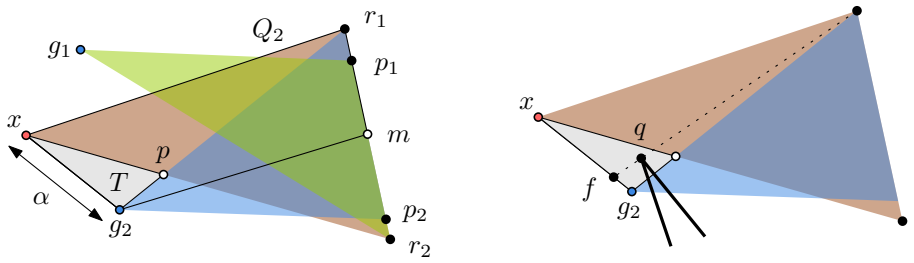

Figure 15: Illustration to the proof of Lemma 14.

At first consider the case that there is a grid point $g \in \alpha\text{-grid}(x)$ inside Δ. This implies the claim immediately, as Δ is convex and it must be empty of obstructions, as Δ can be seen by x.

So, let us assume that there exist two other grid points $g_1, g_2 \in \alpha\text{-grid}(x)$ as in Figure 15. Let us first consider the case that there is no other vertex q with $\text{dist}(x, q) \leqslant L^{-1}$. This implies, we can use Lemma 13 and infer that the visibility regions of g_1, g_2 are only slightly blocked, as indicated in Figure 15. In particular this implies that there exists a point $m \in \text{seg}(r_1, r_2)$ that is visible from g_1 and g_2, as $\text{dist}(r_1, r_2) \geqslant 1 \geqslant 2L^{-2}$. We define the quadrangle $Q_1 = Q(g_1, m, r_2, x)$ and $Q_2 = Q(g_2, m, r_1, x)$. Here $Q(t_1, t_2, t_3, t_4)$ indicates the quadrangle with vertices $t_1, \ldots, t_4$. Clearly it holds $\Delta \subseteq Q_1 \cup Q_2$. Thus it is sufficient to show that g_i sees Q_i, for $i = 1, 2$. We define the point $p = \text{seg}(g_2, r_1) \cap \text{seg}(x, r_2)$ and the triangle $T = \Delta(x, g_2, p)$. Inside T is the only region where part of $\mathcal{P}$ could block the visibility of g_2 to see Q_2 fully. As g_2 can be assumed to see x this blocking part would correspond to a hole of $\mathcal{P}$. But note that a whole has at least 3 vertices and area at least $1/2$. The area of T is bounded by αL, as one of its edges has length α and the height is trivailly bounded by L.

It remains to consider the case that there exists one vertex q with $\text{dist}(x, q) < L^{-1}$. This immediately implies $q \in \alpha\text{-grid}^*(x)$. If q does not block the vision of either g_1 or g_2, we are done. Otherwise, note that q can block the vision of at most one of them, say g_1 and there is *at most one* vertex q with $\text{dist}(x, q) < L^{-1}$. Thus after removal of q the previous argument above can be applied. Now as q blocks g_1 it must be either the case that $q \in \text{cone}(g_1)$ or $q \in T$. In the first case q sees the part of $\text{cone}(g_1)$ that is inside $\mathcal{P}$. In the second case, we denote by f the point $\text{seg}(g_1, x) \cap \ell(r_1, q)$. As $\text{dist}(x, f) \leqslant \text{dist}(x, g_2)$, we can conclude that f and g_1 see Δ by the same argument as above applied to g_1 and f instead of g_1 and g_2, under the assumption that q would not block f. Now however holds that $q \in \text{cone}(f)$ and thus sees what f would have seen and we are done. □

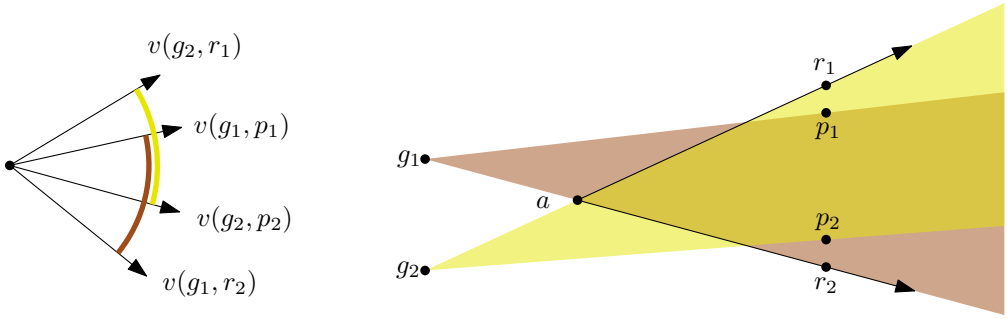

Figure 16: Illustration to Definition 15.

Definition 15 (Cone-Property)**.** Given points g_1, g_2, two reflex vertices r_1 and r_2 and two points $p_1, p_2 \in \text{seg}(r_1, r_2)$ with $\text{dist}(p_1, r_1) \leqslant L^{-2}$ and $\text{dist}(p_2, r_2) \leqslant L^{-2}$, see Figure 16. We denote by C_1 the cone with apex g_1 bounded by the rays $\text{ray}(g_1, r_2)$ and $\text{ray}(g_1, p_1)$ and we

denote by C_2 the cone with apex g_2 bounded by the rays $\text{ray}(g_2, r_1)$ and $\text{ray}(g_2, p_2)$. We say two points g_1, g_2 have the *cone-property* with respect to two reflex vertices r_1, r_2 if there exists some ray contained in $C_1 \cap C_2$.

We define $a = \text{apex}(g_1, g_2) = \text{seg}(g_1, r_2) \cap \text{seg}(g_2, r_1)$.

The definition of C_1 and C_2 might seem a little odd, but in spirit of Lemma 13, we see that $C_i \subseteq \text{cone}(g_i)$, for $i = 1, 2$, if the conditions of Lemma 13 are met.

Lemma 16 (New Cone)**.** *Let g_1, g_2 have the cone-property and assume the notation of Definition 15. Then the cone C with apex a bounded by the rays $ray(a, r_1)$ and $ray(a, r_2)$ is contained in $C_1 \cup C_2$.*

Proof. The directions v with $\text{ray}(g_i, v) \subseteq C_i$, denoted by I_i, form an interval in S^1 for $i = 1, 2$. We define I as the set of directions v such that $\text{ray}(a, v) \subseteq C$. It is easy to see that $I = I_1 \cup I_2$, because the cone-property implies that $I_1 \cap I_2 \neq \varnothing$. Note that for all $b \in C_i$ holds $\text{ray}(b, v) \subseteq C_i$, for all $i = 1, 2$ and $v \in I_i$. The fact $a \in C_1 \cap C_2$ implies $C \subseteq C_1 \cup C_2$. □

Lemma 17 (Cone-Property)**.** *Given two points g_1, g_2 with $dist(g_1, g_2) \leqslant d \leqslant \frac{s}{4} \leqslant \frac{1}{4}$ such that g_1 and g_2 are outside of the embiggened s-bad region of the reflex vertices r_1 and r_2. Then g_1 and g_2 have the cone-property.*

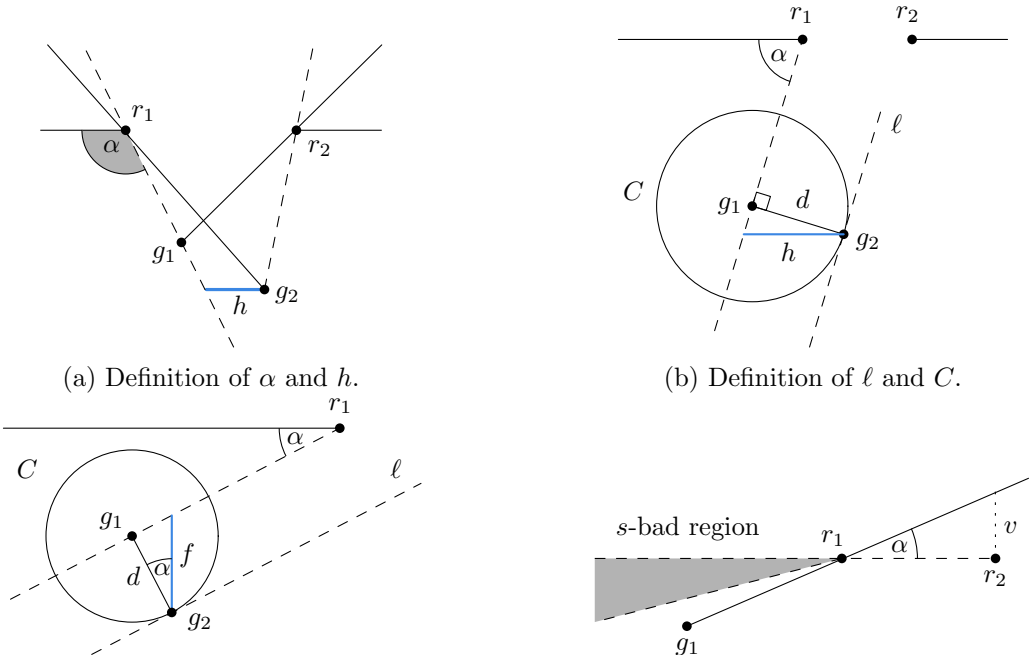

(a) Definition of α and h.

(b) Definition of ℓ and C.

(c) Illustrations for the case $\alpha < 45°$

(d) Illustrations for the case $\alpha < 45°$

Figure 17: Illustrations to the proof of Lemma 17

Proof. It is sufficient to show that the rays $\text{ray}(g_1, p_1)$ and $\text{ray}(g_2, p_2)$ will not intersect, this is they are either parallel or diverge from one another.

Note that $\text{dist}(p_1, p_2) > \text{dist}(r_1, r_2) - \text{dist}(r_1, p_1) - \text{dist}(r_2, p_2) \geqslant 1 - L^{-2} - L^{-2} > 1/2$. Here, we used the fact that all vertices have integer coordinates. If we move r_1 towards p_1 and r_2 towards p_2 then this does not change whether $\text{ray}(g_1, p_1)$ and $\text{ray}(g_2, p_2)$ will intersect or not. Now, the assumption that g_1 and g_2 are not contained in the *embiggened* s-bad region, becomes just that g_1 and g_2 are not contained in the (ordinary) s-bad region.

Thus from now on, we assume $r_1 = p_1$ and $r_2 = p_2$. We cannot make the assumption anymore that r_1 and r_2 have integer coordinates, but we can assume that $\text{dist}(r_1, r_2) \geqslant 1/2$, which is sufficient for the rest of the proof.

Without loss of generality, we assume that the supporting line of the two reflex vertices is horizontal. We can assume that the distance $\text{dist}(g_1, g_2) = d$, as this is the worst case. Further, we assume that g_1 is closer to r_1 than to r_2.

For this purpose, we distinguish two different cases. Either the angle α between $\ell(r_1, r_2)$ and the ray $\text{ray}(g_1, r_1)$ is $\geqslant 45°$ or $< 45°$ degree, see Figure 17a.

In the first case, we compare the *horizontal* distance between the lines at two different locations. To be precise, we will show

$$\text{dist}_{\text{horizontal}}(\ell(g_1, r_1), g_2) < \text{dist}_{\text{horizontal}}(\ell(g_1, r_1), r_2).$$

This shows that the rays are in fact diverging. The horizontal distance between r_2 and $\ell(g_1, r_1)$ equals $\text{dist}(r_1, r_2) \geqslant 1/2$ as was remarked above. The horizontal distance h between the $\ell(g_1, r_1)$ and g_2 can be upper bounded by $\sqrt{2}d < 2d$ as follows. See Figure 17b for an illustration. At first let $\overline{\ell}$ be a line parallel to $\ell(g_1, r_1)$ containing g_2. Then any point p on $\overline{\ell}$ has the same horizontal distance to $\ell(g_1, r_1)$. Further g_2 has distance d to g_1 and thus lies on the circle $C = \partial\text{disk}(g_1, d)$ indicated in Figure 17b. The line ℓ parallel to $\ell(g_1, r_1)$ and furthest away from it that is still intersecting C is indicated in Figure 17b. We can assume that g_2 lies on the intersection of C and ℓ, as h would be smaller in any other case. We draw the horizontal segment t realizing the horizontal distance between g_2 and $\ell(g_1, r_2)$. Note that the angle between $\ell(g_1, r_2)$ and t equals α. It is easy to see that $\text{seg}(g_1, g_2)$ and ℓ are orthogonal. This implies that $\text{seg}(g_1, g_2)$ and $\ell(g_1, r_1)$ are orthogonal as well. It follows

$$\sin\alpha = \frac{d}{h} \quad \Rightarrow \quad h = \frac{d}{\sin\alpha} \leqslant \sqrt{2}d < 2d.$$

Here we used the fact that $\sin\alpha \geqslant \frac{1}{\sqrt{2}}$, for $\alpha \geqslant 45°$. In summary we have

$$\text{dist}_{\text{horizontal}}(\ell(g_1, r_1), g_2) \leqslant 2d \leqslant s/2 \leqslant 1/2 \leqslant \text{dist}_{\text{horizontal}}(\ell(g_1, r_1), r_2).$$

In the second case, we compare the *vertical* distance v between $\ell(g_1, r_1)$ and r_2 and $\ell(g_2, r_2)$ and g_1. We will show

$$\text{dist}_{\text{vertical}}(\ell(g_1, r_1), g_2) \leqslant \text{dist}_{\text{vertical}}(\ell(g_1, r_1), r_2).$$

By the same argument as in case one, we can conclude that

$$f = \text{dist}_{\text{vertical}}(\ell(g_1, r_1), g_2) \leqslant \sqrt{2}d \leqslant 2d.$$

We repeat the argument to avoid potential confusion, see Figure 17c. Let $\overline{\ell}$ be a line parallel to $\ell(g_1, r_1)$. Then every point p on $\overline{\ell}$ has the same vertical distance to $\ell(g_1, r_1)$. Further g_2 has distance d to g_1 and thus lies on the circle $C = \partial\text{disk}(g_1, d)$ indicated in Figure 17c. The line ℓ parallel to $\ell(g_1, r_1)$ and furthest away from it that is still intersecting C is indicated in Figure 17c. We can assume that g_2 lies on the intersection of C and ℓ, as f would be smaller in any other case. We draw the vertical segment t realizing the vertical distance between g_2 and $\ell(g_1, r_2)$. Note that the angle between ℓ and t equals α. It is easy to see that $\text{seg}(g_1, g_2)$ and ℓ are orthogonal. This implies that $\text{seg}(g_1, g_2)$ and $\ell(g_1, r_1)$ are orthogonal as well. It follows

$$\cos\alpha = \frac{d}{f} \quad \Rightarrow \quad f = \frac{d}{\cos\alpha} \leqslant \sqrt{2}d < 2d.$$

Here we used the fact that $\cos\alpha \geqslant \frac{1}{\sqrt{2}}$, for $\alpha < 45°$.

Now, we give some bounds on the vertical distance v between $\ell(g_1, r_1)$ and r_2. See Figure 17d. Note that $\tan\alpha \geqslant s$ by the assumption that g_1 is not contained in an s-bad region. This implies

$$s \leqslant \tan\alpha = \frac{v}{\operatorname{dist}(r_1, r_2)} \leqslant \frac{v}{1/2} \leqslant 2v.$$

In summary we have

$$\operatorname{dist}_{\text{vertical}}(\ell(g_1, r_1), g_2) \leqslant 2d \overset{(1)}{\leqslant} s/2 \leqslant v = \operatorname{dist}_{\text{vertical}}(\ell(g_1, r_1), r_2).$$

For (1) we used the assumption of the lemma. □

Definition 18 (Cone Containment and Cutting Cones)**.** We say $\operatorname{cone}(x)$ is contained in $\operatorname{cone}(y)$ *behind* r_1 and r_2 if $\operatorname{cone}(x) \cap \ell^+ \subseteq \operatorname{cone}(y) \cap \ell^+$, where ℓ^+ is the half-plane bounded by $\ell(r_1, r_2)$ and does not contain the points x and y. When r_1 and r_2 are clear from context, we just say $\operatorname{cone}(x)$ is contained in $\operatorname{cone}(y)$. In the same fashion, we define $\operatorname{cone}(z) = \operatorname{cone}(x) \cup \operatorname{cone}(y)$ *behind* r_1 and r_2. We say some cone C is *cut* by a line segment s if the line segment $s' = C \cap s$ is non-empty and contains neither end point of s.

It is easy to see that for any $\operatorname{cone}(x)$ either holds that there exists a point $g \in \alpha\text{-grid}(x)$ with $g \in \operatorname{cone}(x)$ or there exists two points $g_1, g_2 \in \alpha\text{-grid}(x)$ such that $\operatorname{cone}(x)$ cuts $\operatorname{seg}(g_1, g_2)$, see Figure 11.

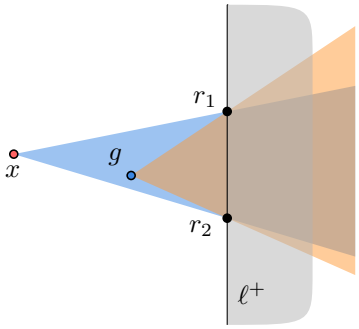

(a) The point g is contained in $\operatorname{cone}(x)$ and thus $\operatorname{cone}(g)$ contains $\operatorname{cone}(x)$ behind ℓ^+.

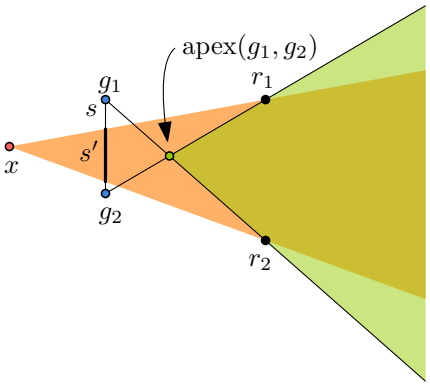

(b) The line segment s cuts $\operatorname{cone}(x)$, as s' is non-empty and contains neither endpoint of s. Thus $\operatorname{cone}(x)$ contains $\operatorname{apex}(x, y)$.

Figure 18: Illustrations to cone containment and cone cutting.

Lemma 19 (Cone-Containment)**.** *Consider the cones of x and g with respect to r_1 and r_2. Let $g \in \operatorname{cone}(x)$ then $\operatorname{cone}(x)$ is contained $\operatorname{cone}(g)$ behind r_1 and r_2.*

Lemma 20 (Cut-Segments)**.** *Consider the cones of g_1, g_2 and x with respect to r_1 and r_2. Let g_1, g_2 have the cone-property and assume that $\operatorname{cone}(x)$ is cut by $\operatorname{seg}(g_1, g_2)$. Then $\operatorname{cone}(x)$ is contained in $\operatorname{cone}(a)$, where $a = \operatorname{apex}(g_1, g_2)$ behind r_1 and r_2. Further $\operatorname{cone}(a) = \operatorname{cone}(g_1) \cup \operatorname{cone}(g_2)$ behind r_1 and r_2.*

Proof. We will show that $a \in \operatorname{cone}(x)$. See Figure 18b for an illustration of the proof. Let h_1 be the half-plane bounded by $\ell(x, r_1)$ containing r_2 and let h_2 be the half-plane bounded by $\ell(x, r_2)$ containing r_1. Recall that $a = \operatorname{apex}(g_1, g_2) = \operatorname{seg}(g_2, r_1) \cap \operatorname{seg}(g_1, r_2)$. Then it is clear that $a \in \operatorname{seg}(g_2, r_1) \subseteq h_1$ and $a \in \operatorname{seg}(g_1, r_2) \subseteq h_2$. Thus $a \in h_1 \cap h_2 = \operatorname{cone}(x)$. □

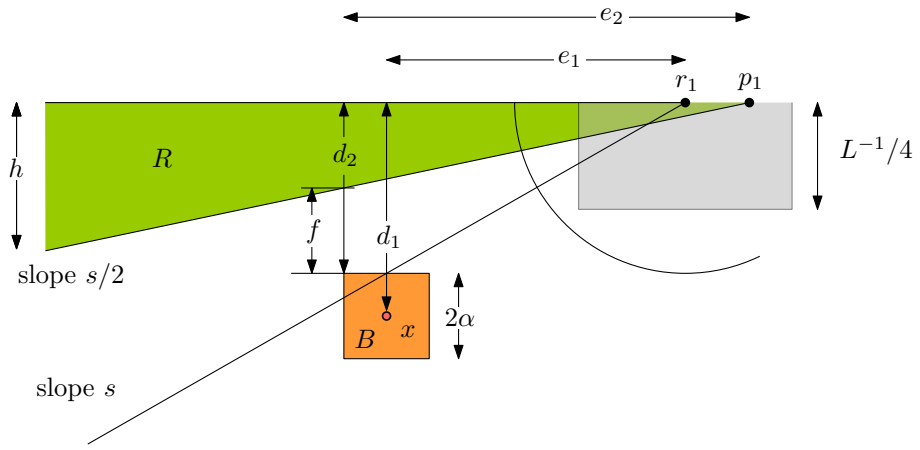

Figure 19: Illustrations to Lemma 21. $\ell = \ell(r_1, r_2)$ is horizontal; B is an axis parallel box around x with side length 2α; R is an embiggened $(s/2)$-bad region; e_1 is the horizontal distance between x and r_1; $e_2 = e_1 + \alpha + L^{-2}$; d_1 is the vertical distance between ℓ and x; d_2 is the vertical distance between ℓ and B; $f = d_2 - (s/2)e_2$.

Lemma 21 (Grid Outside Bad Region)**.** *Let x be a point not in the s-bad region of r_1, r_2, seeing r_1 and r_2, and dist$(x, r_i) \geqslant L^{-1}$, for $i = 1, 2$. Further we assume $s \leqslant 8L\alpha$ and $L^{-3} \leqslant s$. Then it holds that α-grid(x) is not in the embiggened $\frac{s}{2}$-bad region of r_1 and r_2.*

Proof. Refer to Figure 19 for an illustration of this proof and the notation therein. We assume that $\ell(r_1, r_2)$ is horizontal and x is closer to r_1 then to r_2. As every point of α-grid(x) has distance at most α from x, it is sufficient to show that the box B with sidelength 2α centered at x does not intersect the embiggened $\frac{s}{2}$-bad region R. For the proof it is sufficient to assume that x is outside the box C of sidelength $L^{-1}/2$ centered at r_1.

We denote by h the largest width of R. It is easy to see that $h \leqslant (s/2)L \leqslant L^{-2}$.

In case that $e_1 < L^{-1}/2$ holds $d_2 = d_1 - \alpha > L^{-1}/4 \gg h$. Thus, we assume from now on $e_1 \geqslant L^{-1}/2$. In this case, it holds

$$e_2 = e_1 + L^{-2} + \alpha < \tfrac{3e_1}{2}.$$

It suffices to show $f > 0$.

$$f = d_2 - \frac{e_2 s}{2} \overset{(1)}{>} d_2 - \frac{3e_1 s}{4} \overset{(2)}{=} d_1 - \alpha - \frac{3e_1 s}{4} \overset{(3)}{\geqslant} e_1 s - \alpha - \frac{3e_1 s}{4} = \tfrac{e_1 s}{4} - \alpha \overset{(4)}{\geqslant} \frac{L^{-1} s}{8} - \alpha \overset{(5)}{>} 0.$$

Inequality (1) was just shown above. Equality (2) can be easily seen in Figure 19. Inequality (3) applies the fact that the s-bad region does not contain x and has slope s. Inequality (4) follows the assumed lower bound $e_1 > L^{-1}/4$. Inequality (5) is the main assumption of the lemma. □

We are now ready to proof Lemma 11.

Lemma 11 (Special Local Visibility Containment Property)**.** *Let r_1 and r_2 be two consecutive vertices in the clockwise order of the vertices visible from $x \in \mathcal{P}$ and let x be outside the s-bad region of the vertices r_1 and r_2. We make the following assumptions:*

1. $s \leqslant L^{-3}$
2. $\alpha \leqslant L^{-7}$.
3. $16L\alpha \leqslant s$

Then α-grid$^*(x)$ sees cone(x).

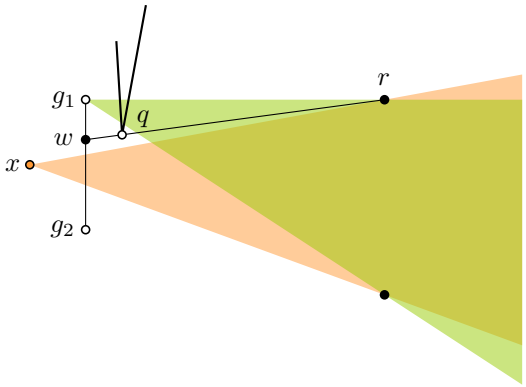

Figure 20: In this scenario there exists a vertex q, which blocks the visibility g_1. Luckily, we can define another point w, which is not blocked by q. And q can see the cone of w.

Proof of Lemma 11. Note first that the triangle $\Delta(x, r_1, r_2)$ is seen by α-grid$^*(x)$ by Lemma (Small Triangle) 14. So, we are only interested in cone(x) behind r_1 and r_2.

The easiest case to be ruled out is that there exists $g \in \alpha$-grid$^*(x)$ such that $g \in \text{cone}(x)$, as by Lemma (Cone-Containment) 19 this implies the claim. So from now on, we assume for all $g \in \alpha$-grid$^*(x)$ holds $g \notin \text{cone}(x)$.

We know that there exists two grid points g_1 and g_2 such that cone(x) cuts the segment seg(g_1, g_2). By Lemma (Cut-Segments) 20 it remains to show that g_1 and g_2 have the Cone-Property. For this purpose, we want to invoke Lemma (Cone-Property) 17. To this end, we have to invoke a series of other lemmas.

Note that $r_1, r_2 \in \text{cone}(x)$ and $\text{dist}(x, r_i) < L^{-1}$, for $i = 1$ or $i = 2$ implies r_i is included in α-grid$^*(x)$. Thus by the argument above, we assume from now on $\text{dist}(x, r_i) \geqslant L^{-1}$ for all $i = 1, 2$. There might still be a different vertex $q \neq r_1, r_2$ with $\text{dist}(x, q) < L^{-1}$. We deal first with the case that there are no vertices q with $\text{dist}(x, q) < L^{-1}$. By Lemma (Grid Outside Bad Region) 21 holds that α-grid(x) is not contained in the embiggened s'-bad region with $s' = s/2$ with respect to r_1 and r_2. We can apply Lemma (Limited Blocking) 13 as $\alpha < L^{-7}$ and what is said above. From Lemma (Limited Blocking) 13 follows that cone(g_i) contains C_i, for $i = 1, 2$ as defined in Definition (Cone-Property) 15, see also Definition (Cone-Property) 10 to recall the definition of cone(g_i).

It remains to show that g_1 and g_2 satisfy the Cone-Property. To this end, we need to show that the assumptions of Lemma (Cone-Property) 17 is met. Here, we consider the embiggened s'-bad region with $s' = s/2$. Note that

$$\text{dist}(g_1, g_2) \leqslant 2\alpha \leqslant s/(8L) < s'/4 = s/8 < 1/4.$$

This shows the claim together with Lemma (New Cone) 16.

It remains to consider the case that there exists one vertex q with $\text{dist}(x, q) < L^{-1}$, see Figure 20. This immediately implies $q \in \alpha$-grid$^*(x)$. If q does not block the vision of either g_1 or g_2, we are done. Otherwise, note that q can block the vision of at most one of them, say g_1 and there is *at most one* vertex q with $\text{dist}(x, q) < L^{-1}$. We define $w = \ell(q, r) \cap \text{seg}(g_1, g_2)$. As the edges incident to q block g_1 at least partially, we know that w exists. As $\text{dist}(w, g_2) \leqslant \text{dist}(g_1, g_2)$ and cone(x) cuts the segment seg(w, g_2), we can use the same arguments as above. By definition of w holds $q \in \text{cone}(w)$. And behind r_1, r_2 the cone(q) contains cone(w). □

4.5 Global Visibility Containment

This simple lemma quantifies (as a function of s and L) the maximum *width* of s-bad regions.

Lemma 22 (distance bad region to supporting line)**.** *Let p be a point of $\mathcal{P}$ inside an s-bad region associated to opposite reflex vertices r_1 and r_2. Then dist$(p, \ell(r_1, r_2)) \leqslant sL$.*

Proof. See Figure 12a. □

Although it is not possible to achieve a local visibility containment property for all points in $\mathcal{P}$, the exceptions only involve bad regions. Under Assumption 2 (and Assumption 1), we can give a fairly short proof of Lemma 2. As preparation, we need the following technical lemma which heavily relies on Lemma 5.

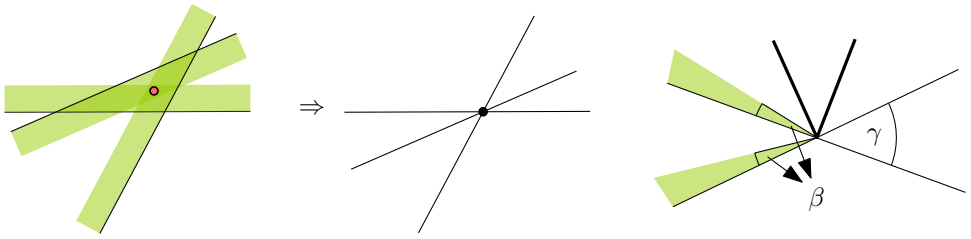

Figure 21: Three bad regions meeting in an interior point implies that the extensions must meet in a single point. No two bad regions intersect in the vicinity of a vertex, as they are defined by some angle $\beta \ll L^{-2}$. But the angle γ between any two extensions is at least L^{-2}.

Lemma 23 (no three bad regions intersect)**.** *Under Assumption 1 and 2, for any $s \leqslant L^{-9}$, no point in the interior of $\mathcal{P}$ belongs to three different s-bad regions.*

Proof. We consider now the case that there exists a point x with $\mathrm{dist}(x, v) \leqslant L^{-2}$, for some vertex v. We show that x is contained in at most *one* bad regions, see to the right of Figure 21. note first that any extension ℓ with $v \notin \ell$ has distance at least L^{-1} from v, by Lemma 5. And thus x cannot be in any bad region belonging to ℓ by Lemma 22. By Lemma 5 the angle between any two extensions must be at least L^{-2}. As we have at most two vertices contained on any line, the bad regions belonging to the extensions through v must start at v, see Definition 9. For the angle β defining the bad regions holds $\tan(\beta) = s \ll L^{-2}$ and thus all bad regions in the vicinity of v are disjoint. (Note that v itself is not considered as part of the bad region.)

Let ℓ_1, ℓ_2, ℓ_3 be supporting lines of three distinct pairs of opposite reflex vertices. As we assumed that no three points lie on a line, those three lines are also distinct. We first consider the case where two of those supporting lines, say ℓ_1 and ℓ_2 are parallel. By Lemma 5 Item 5, $\mathrm{dist}(\ell_1, \ell_2) \geqslant L^{-1}$. Also, by Lemma 22, any point of an s-bad region is at distance at most sL of the corresponding supporting line. Therefore, any point in the intersection of the s-bad region associated to ℓ_1 and the one associated to ℓ_2 is at distance at most L^{-8} from those two lines; a contradiction to $\mathrm{dist}(\ell_1, \ell_2) \geqslant L^{-1}$.

We now show that any intersection of two supporting lines (among ℓ_1, ℓ_2, ℓ_3) should be in the interior of $\mathcal{P}$. Such an intersection cannot be on the boundary of $\mathcal{P}$ deprived of the vertices of $\mathcal{P}$, since it would immediately yield three supporting lines meeting in a point. If two supporting lines, say ℓ_1 and ℓ_2, meet in a vertex of $\mathcal{P}$, then this vertex is one of the opposite reflex vertices for both ℓ_1 and ℓ_2 (otherwise there would be three vertices on a line).

Assume now that the intersection p of say, ℓ_1 and ℓ_2 is outside $\mathcal{P}$. By Lemma 5 Item 4, the distance of p to any point in $\mathcal{P}$ is at least L^{-5}. Let p' be a point of $\mathcal{P}$ in the intersection of two s-bad regions associated to ℓ_1 and to ℓ_2. By Lemma 22, the distance of p' to both ℓ_1 and ℓ_2 is at most L^{-8}. That implies, by setting d to L^{-8} in Lemma 5 Item 7, that $\mathrm{dist}(p', p) \leqslant L^{-8}L^2 = L^{-6}$; a contradiction to $\mathrm{dist}(p, \mathcal{P}) \geqslant L^{-5}$

Thus, we can suppose that ℓ_1, ℓ_2, ℓ_3 pairwise intersect in three distinct points $p = \ell_1 \cap \ell_2, q = \ell_1 \cap \ell_3, r = \ell_2 \cap \ell_3$ in the interior of $\mathcal{P}$; this is because we assumed that no three extensions meet in a point. Let p' be in the s-bad regions associated to ℓ_1 and ℓ_2. As explained in the end of the previous paragraph, $\mathrm{dist}(p', p) \leqslant L^{-6}$. By Lemma 5 Item 3, $\mathrm{dist}(p, \ell_3) \geqslant L^{-5}$. By Lemma 22 any point in the s-bad region associated to ℓ_3 is at distance at most L^{-8}. As $L^{-6} + L^{-8} < L^{-5}$,

p' can not be in the s-bad region associated to ℓ_3. Which means that the intersection of the three s-bad regions associated to ℓ_1, ℓ_2, and ℓ_3 is empty. □

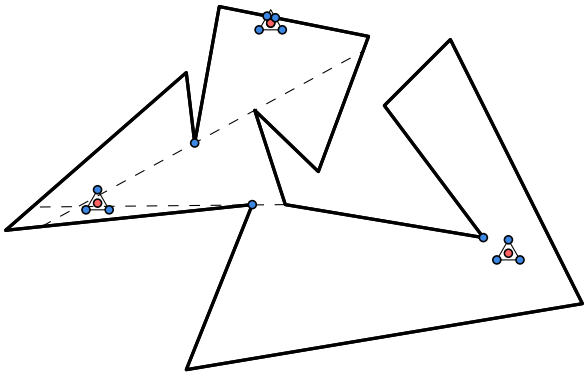

Figure 22: The red dots indicate the optimal solution. The blue dots indicate the The red dot on the top is in the interior case and four grid points are added around it. The red dot on the left is too close to two supporting lines and we add one of the reflex vertices of each of the supporting lines. The red dot to the right has distance less than L^{-1} to a reflex vertex, so we add that vertex as well to G.

Proof Lemma 2 using Assumptions 1 and 2. We denote by OPT an optimal solution of size k. We assume that no point of OPT is actually contained in Γ as we can just take that point into our grid solution. In particular this implies OPT contains non of the vertices of $\mathcal{P}$. Let $\alpha = L^{-11}$ and $s < L^{-9}$. Let $x \in \mathcal{P}$ be some point and $R(x)$ some set of size at most 2 that contains a reflex vertex for each s-bad region, that x is contained in. As no point is contained in three bad regions $R(x)$ having size 2 is enough, see Lemma 23. We define

$$G = \bigcup_{x \in OPT} \alpha\text{-grid}^*(x) \cup R(x).$$

It is easy to see that $G \subseteq \Gamma$ has size $|G| \leqslant (7+2)k$. We want to argue that G sees the entire polygon. For each $x \in OPT$ the local containment property holds, except for the bad regions it is in, see Lemma 11. These parts are seen by the reflex vertices we added. □

5 Conclusion

We presented an $O(\log |OPT|)$-approximation algorithm for the Point Guard Art Gallery problem under two relatively mild assumptions. The most natural open question is whether Assumption 2 can be removed. We believe that this is possible but it will require some additional efforts and ideas. Another improvement of the result would be to achieve an approximation ratio of $O(\log n)$ for polygons with holes. This would match the currently best known algorithm for the Vertex Guard variant and known lower bounds. In that respect, it might be very useful that Lemma 2 does not require the polygon to be simple. One might also ask about the inapproximability of Point Guard Art Gallery for simple polygons. For the moment, the problem is only known to be inapproximable for a certain constant ratio (quite close to 1), unless P=NP. It would be interesting to get superconstant inapproximability under standard complexity theoretic assumptions or improved approximation algorithms.

References

[1] E. S. Ayaz and A. Üngör. Minimal witness sets for art gallery problems. *EuroCG*, 2016.

[2] J. Barát, V. Dujmovic, G. Joret, M. S. Payne, L. Scharf, D. Schymura, P. Valtr, and D. R. Wood. Empty pentagons in point sets with collinearities. *SIAM J. Discrete Math.*, 29(1):198–209, 2015.

[3] S. Basu, R. Pollack, and M.-F. Roy. Algorithms in real algebraic geometry. *AMC*, 10:12, 2011.

[4] P. Belleville. Computing two-covers of simple polygons. Master's thesis, McGill University, 1991.

[5] V. V. S. P. Bhattiprolu and S. Har-Peled. Separating a voronoi diagram via local search. In *SOCG 2016*, pages 18:1–18:16.

[6] É. Bonnet and T. Miltzow. The parameterized hardness of the art gallery problem. In *ESA 2016*, page to appear, 2016.

[7] É. Bonnet and T. Miltzow. The parameterized hardness of the art gallery problem. *CoRR*, abs/1603.08116, 2016.

[8] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[9] J. Canny. Some algebraic and geometric computations in pspace. In *STOC*, pages 460–467. ACM, 1988.

[10] J. Cardinal. Computational geometry column 62. *SIGACT News*, 46(4):69–78, Dec. 2015.

[11] K. Chwa, B. Jo, C. Knauer, E. Moet, R. van Oostrum, and C. Shin. Guarding art galleries by guarding witnesses. *Int. J. Comput. Geometry Appl.*, 16(2-3):205–226, 2006.

[12] K. L. Clarkson. Algorithms for polytope covering and approximation. In *WADS 1993*, pages 246–252, 1993.

[13] A. Deshpande. A pseudo-polynomial time $O(\log^2 n)$-approximation algorithm for art gallery problems. Master's thesis, Department of Mechanical Engineering, Department of Electrical Engineering and Computer Science, MIT, 2006.

[14] A. Deshpande, T. Kim, E. D. Demaine, and S. E. Sarma. A pseudopolynomial time $O(\log n)$-approximation algorithm for art gallery problems. In *WADS 2007*, pages 163–174, 2007.

[15] S. Durocher and S. Mehrabi. Guarding orthogonal art galleries using sliding cameras: algorithmic and hardness results. In *MFCS 2013*, pages 314–324. Springer, 2013.

[16] A. Efrat and S. Har-Peled. Guarding galleries and terrains. *Inf. Process. Lett.*, 100(6):238–245, 2006.

[17] S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.

[18] S. Fisk. A short proof of chvátal's watchman theorem. *J. Comb. Theory, Ser. B*, 24(3):374, 1978.

[19] S. K. Ghosh. *Visibility algorithms in the plane*. Cambridge University Press, 2007.

[20] S. K. Ghosh. Approximation algorithms for art gallery problems in polygons. *Discrete Applied Mathematics*, 158(6):718–722, 2010.

[21] A. Gilbers and R. Klein. A new upper bound for the vc-dimension of visibility regions. *Computational Geometry*, 47(1):61–74, 2014.

[22] G. Kalai and J. Matoušek. Guarding galleries where every point sees a large area. *Israel Journal of Mathematics*, 101(1):125–139, 1997.

[23] H. Kim and G. Rote. Congruence testing of point sets in 4 dimensions. *CoRR*, abs/1603.07269, 2016.

[24] H. Kim and G. Rote. Congruence testing of point sets in 4-space. In *Symposium on Computational Geometry*, volume 51 of *LIPIcs*, pages 48:1–48:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.

[25] J. King. Fast vertex guarding for polygons with and without holes. *Comput. Geom.*, 46(3):219–231, 2013.

[26] D. G. Kirkpatrick. An $O(\log \log OPT)$-approximation algorithm for multi-guarding galleries. *Discrete & Computational Geometry*, 53(2):327–343, 2015.

[27] E. A. Krohn and B. J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013.

[28] J. Matousek. Intersection graphs of segments and $\exists\mathbb{R}$. *CoRR*, abs/1406.2636, 2014.

[29] R. Motwani, A. Raghunathan, and H. Saran. Covering orthogonal polygons with star polygons: The perfect graph approach. *J. Comput. Syst. Sci.*, 40(1):19–48, 1990.

[30] J. O'rourke. *Art gallery theorems and algorithms*, volume 57. Oxford University Press Oxford, 1987.

[31] M. Schaefer. *Complexity of Some Geometric and Topological Problems*, pages 334–344. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[32] T. C. Shermer. Recent results in art galleries. *Proceedings of the IEEE*, 80(9):1384–1399, 1992.

[33] J. Urrutia et al. Art gallery and illumination problems. *Handbook of computational geometry*, 1(1):973–1027, 2000.

[34] P. Valtr. Guarding galleries where no point sees a small area. *Israel Journal of Mathematics*, 104(1):1–16, 1998.